



INVESTOR IN PEOPLE


The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated

29 November 2002

The
Patent
Office


Patents Act 1977

24 OCT 2002

Rule 16

25OCT02 E758446-1 000611
P01/7700 0.00-0224806.0

Request for grant of a patent



The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920020053GB1		
2.	Patent application number (The Patent Office will fill in this part)	0224806.0		
3.	Full name, address and postcode of the or of each applicant (underline all surnames)	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America 519637001		
	Patents ADP number (if you know it)			
	If the applicant is a corporate body, give the country/state of its incorporation	State of New York United States of America		
4.	Title of the invention	METHOD AND APPARATUS FOR A INTERACTIVE VOICE RESPONSE SYSTEM		
5.	Name of your agent (if you have one)	P Waldner		
	"Address for Service" in the United Kingdom to which all correspondence should be sent (including the postcode)	IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire S021 2JN 919006		
	Patents ADP number (if you know it)			
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority App No (if you know it)	Date of filing (day/month/year)
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date or the earlier application	No of earlier application	Date of filing (day/month/year)	

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.)
- Yes

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 10

Claim(s) 2

Abstract 1

Drawing(s) 6 *ib* *hm*

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 3

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)


11. I/We request the grant of a patent on the basis of this application

P Waldner
Signature

22 October 2002
Date

12. Name and daytime telephone number of person to contact in the United Kingdom P Waldner
01962 816057

Patents Act 1977
Rule 15



7/77

Statement of inventorship and of right to grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

24 OCT 2002

1.	Your reference	GB920020053GB1
2.	Patent application number (if you know it)	0224806.0
3.	Full name of the or of each applicant	INTERNATIONAL BUSINESS MACHINES CORPORATION
4.	Title of invention	METHOD AND APPARATUS FOR A INTERACTIVE VOICE RESPONSE SYSTEM
5.	State how the applicant(s) derived the right from the inventor(s) to be granted a patent	By employment and by agreement
6.	How many, if any, additional Patents Forms 7/77 are attached to this form?	
7.	I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.	
	Signature	22 October 2002 Date
8.	Name and daytime telephone number of person to contact in the United Kingdom	P Waldner Tel: 01962 815830

Patents Form 7/77

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Wendy-Ann COYLE
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

8492233001

Stephen James HASKEY
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

8492241001

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

REMINDER

Have you signed the form?

Patents ADP number (if known)

METHOD AND APPARATUS FOR A INTERACTIVE VOICE RESPONSE SYSTEM

This invention relates to a method and apparatus for an interactive voice response system.

BACKGROUND

A telephone can be used to place a catalogue order; check an airline schedule; query a price; review an account balance; notify a customer; record and retrieve a message; and many other business interactions. Often, each telephone call involves a service representative talking to a user, asking questions, entering responses into a computer, and reading information to the user from a terminal screen. This process can be automated by substituting an interactive voice response system (IVR) with an ability to play voice prompts and receive user input e.g. from speech recognition or from DTMF tones.

The interaction of the voice prompts and user input is guided by a voice application that in turn is executed by the IVR. Voice applications have been written in script, state code, Java*, and voice extensible markup language (VoiceXML). *Java and all Java based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc in the United States, other countries or both.

When a known user speaks correctly, clearly and in a good quiet environment, an IVR speech recognition engine trained in the speech of the user has little problem understanding the speech utterances. However, if the speech differs too much from the data that was used to train the speech recognition engine then understanding the speech can be difficult. When the system cannot correctly recognize an utterance it is usual for the engine to prompt the user to repeat the utterance. However, if the user is unaware of how to help the IVR he can compound the problem. For instance it may be that the user is speaking too slowly in the first place and speaking even more slowly will compound the problem. This is because it is human nature to speak louder, slower or in an over pronounced way when not initially understood. This can often be heard when a tourist abroad tries to speak, in their native tongue, to a foreigner. This may help when speaking to a human but usually has the opposite effect when speaking to a machine.

US patent publication 6314396 discloses a system for automatic gain adjustment in a speech recognition system and returns a signal to noise ratio to a user with the intention of achieving a better speech recognition accuracy.

Therefore, although it is known to give feedback to a user as to the signal to noise ratio, no-one has disclosed feedback concerning the speed at which a user is speaking.

SUMMARY OF INVENTION

According to a first aspect of the present invention there is provided an interactive voice response system comprising: an prompt acquisition component for acquiring an utterance from a user; a speech recognition engine for recognising a plurality of words from the utterance; a custom server for comparing the actual duration of the utterance with an ideal duration of the recognised words; and a prompt play component for prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

In this way, data available from a speech recognition engine is used to estimate the speed at which the user is speaking by comparing an ideal duration of the recognised words (as stored with the model of the speech data in the speech recognition engine) with the actual duration of the spoken words.

Preferably the means for comparing the actual duration of the utterance with an ideal duration of the recognised words comprise means for acquiring for each word the actual duration of delivery and ideal duration and means for comparing the differences in actual duration and ideal duration for each word. This solution breaks an utterance down into component words and calculates a difference for each word and then finds the average of all the words. The advantage is that for each recognised word there already exists an ideal duration value in the speech model. The means for acquiring and the means for comparing are defined in the duration custom server.

Advantageously, the means for comparing the difference between actual duration and ideal duration of each word comprises: calculating the ratio of actual duration and ideal duration for each word wherein the ratio is an indication of the speed of delivery of the utterance. Using a ratio allows relative comparisons rather than an absolute comparison which would compare small words differently than large words.

More advantageously the means for comparing the actual duration of the utterance with an ideal duration of the recognised words comprises calculating an average of the ratio of words as an indication of the speed of delivery of the utterance. Such an average allows a comparison to view

the whole picture rather than individual ratios which may on their own distort any conclusion.

Advantageously, each word has an associated recognition score and the ratio of a word is only used in the average calculation if the associated recognition score is above a lower threshold recognition score. As the overall recognition score for a utterance diminishes then the reliability having the correct word with corresponding duration also diminishes.

If the total actual duration is greater than the total ideal duration then a prompt that the user is speaking too fast might help the next acquisition. If the total actual duration is less than the total ideal duration then a prompt that the user is speaking too slow might likewise help.

Most preferably the means for prompting the user only operates if the actual duration and the ideal duration differ by more than a de minimus value. This solution bypasses the prompt for feedback when there are few problems with the recognition. If the recognition is perfect then the user will become annoyed at hearing prompts to speak faster or slower.

Advantageously, the system further comprises means for re-acquiring the utterance from the user in case the original utterance cannot be identified.

According to a second aspect of the invention there is provided a method in an interactive response system comprising: acquiring an utterance from a user; recognising a plurality of words from the utterance; comparing the actual duration of the utterance with an ideal duration of the recognised words; and prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

According to a third aspect of the invention there is provided a computer program product for processing one or more sets of data processing tasks, said computer program product comprising computer program instructions stored on a computer-readable storage medium for, when loaded into a computer and executed, causing a computer to carry out the steps of: acquiring an utterance from a user; recognising a plurality of words from the utterance; comparing the actual duration of the utterance with an ideal duration of the recognised words; and prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

DESCRIPTION OF DRAWINGS

In order to promote a fuller understanding of this and other aspects of the present invention, an embodiment of the invention will now be described, by means of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic of an voice telephony system;

Figure 2 is a more detailed schematic of the IVR system of the preferred embodiment of the present invention;

Figure 3 is a schematic application process of the preferred embodiment of the present invention;

Figure 4 is a schematic custom server process of the preferred embodiment of the present invention;

Figure 5A and 5B are example utterances as stored in an utterance database according to a preferred embodiment of the present invention; and

Figure 6 is a table of data received from a speech recognition engine according the preferred embodiment of the present invention.

DESCRIPTION OF THE EMBODIMENTS

According to Figure 1 there is shown a schematic of a voice telephony in the which the present invention can be embodied. Voice telephony system 100 comprises an interactive voice response system (IVR) 102 connected to a voice server 116 over an LAN 114. An example of IVR 102 is *IBM *WebSphere Voice Response 3.1 (WVR) for AIX based on IBM *DirectTalk Technology 102. An example of voice server 116 is IBM Voice Server. A user uses a telephone 106 to connect with IVR 102 through telephony (PSTN) switch (PABX) 104. IVR 102 uses any one of its three application languages to control a voice interaction. Java application layer 108 uses Java Beans and Java application to control the IVR 102. State table environment 110 hosts the original DirectTalk application programming language and is based on state table applications and custom servers. VoiceXML application layer 112 uses VoiceXML browsers and VoiceXML applications in Web Servers to control the IVR 100.

IVR 102 is well-suited for large enterprises or telecommunications businesses. It is scalable, robust and designed for continuous operation 24 hours a day and 7 days a week. IBM WebSphere Voice Response 3.1 for AIX

can support between 12 and 480 concurrent telephone channels on a single system. Multiple systems can be networked together to provide larger configurations. *AIX, DirectTalk, IBM, pSeries, and WebSphere are trademarks of International Business Machines in the United States, other countries, or both.

The preferred embodiment uses WebSphere Voice Resonse for AIX 3.1 which supports from 1 to 16 E1 or T1 digital trunks on a single IBM pSeries* server with up to 1,500 ports on a single system. Up to 2304 telephony channels using T1 connections or 2880 telephony channels using E1 connections can be supported in a 19" rack. WebSphere Voice Response for AIX 3.1 requires an IBM AIX v 4.3 operating system running on an IBM pSeries computer. It supports network connectivity on multiple networks including PSTN, ISDN, CAS, SS7, VoIP networks. The preferred embodiment is concerned with those networks which provide a user identification number with an incoming call e.g. ISDN and SS7.

Voice server 116 provides voice processing capacity to the telephony base of IVR 102 in the form of speech recognition engine 118 and text-to-speech engine 120.

The speech recognition engine 118 analyzes input audio using individual pronunciation models for all words in an active vocabulary, including a word representing <silence>. The engine analyzes the audio by fitting it to a mathematical pronunciation model of words in all possible word sequences specified as possible by the vocabulary's language model. The fitting process includes computing a distribution for when each word begins and ends, with the most probable transition points of each distribution reported as the word boundaries. The quality of the mathematical fit between word models and input audio is used together with the language model probability for each word in a particular sequence and several other parameters of the decoding process to compute word scores. During runtime the engine creates these metrics (start time, end time and score) for every word which are passed to the voice custom server along with the recognised word result. An ideal duration time metric of the recognised word result is based on the speech recognition language model. Each phoneme in the language model has an associated ideal duration time and the duration time for a recognised word is the sum of the durations for the phonemes in the recognised word. Normally only the recognised word result is sent to the IVR but the other metrics are available on demand. In the present embodiment the IVR requests all the above metrics with each utterance it sends to the speech recognition engine.

Referring to Figure 2, state table environment 110 of IVR 102 comprises; a state table application 202; a duration custom server 204; a voice custom server 206 and an utterance database 208. State table application 202 controls a voice interaction on the IVR 102 when a voice channel from a telephone is opened. The state table application 202 performs application method 300 which is described in relation to Figure 3. Voice custom server 206 provides the interface to the speech recognition engine 118 and the text-to-speech engine 120 on voice server 116. Voice custom server 206 places the results of speech recognition into the utterance database 208 after a request from the state table application 202. The results of the speech recognition comprise: the recognised words of the utterance; a recognition score for each word; an actual duration for each word; and an ideal duration for each word as used in the speech recognition model. Utterance database 208 receives the results of speech recognition from voice custom server 206 and further processing is performed on the results by duration custom server 204. An example of the results and further processing is shown in Figures 5A and 5B. Duration custom server 204 acquires the data in utterance database 208 and compares the spoken duration of actual word in an utterance with the ideal duration, this is further described with reference to duration custom server method 400 of Figure 4.

Referring to Figure 3, method 300 performed by state table application 202 is described in more detail. The first step is acquiring an utterance (step 302) from a user connected to the IVR 100 after prompting the user to speak into the telephone. Recognising a word string from utterance (step 304) is performed through the custom voice server 206 using the voice server 116 and speech recognition engine 118. The results of the recognition are placed into the utterance database 208. Step 306 calculates a duration ratio. A comparison of the actual duration of utterance with an ideal duration of recognised words is performed by duration custom server 204 by acquiring the values from the utterance database 208. The duration custom server also calculates an average recognition score for the whole utterance which is computed using an average of the recognition scores for all of the words. In step 308, the state table application 202 prompts the user with "please speak a little faster next time" or "please speak a little slower next time" depending on the duration ratio. In this example a duration ratio of more than one is an indication that the user is speaking slower than the ideal speed. A duration ratio of less than one indicates that the user is speaking faster than the ideal speed. The application then re-acquires the utterance (step 310) if there are words with recognition scores below a lower threshold recognition score, that is below 60%. A lower threshold recognition score is different for each speech recognition engine and configuration of the engine so, by way of example

only, 60% is taken as the lower threshold recognition score to explain the embodiment. If there are words in the utterance database with recognition scores below 60% then the application re-acquires the utterance at step 314. Otherwise the method finishes at step 312 and continues with the remainder of the state table application accepting or rejecting the recognised words. In normal operation re-acquisition is only performed once or twice and the best result is used or the result is negated. Step 316 skips prompt step 308 and the re-acquire step 310 if there is no need to prompt the user to speak slower or quicker. This situation occurs when the duration ratio is within a de minimus value, for example between 1.2 and 0.80 but also when the overall recognition value is above an upper threshold recognition score, for example 90%.

Referring to Figure 4, method 400 of the duration custom server 204 is described. The state table application 202 calls the duration custom server 204 after an utterance has been recorded to compare actual duration of the utterance with an ideal duration of the recognised words (step 402). The actual duration in seconds is acquired for the first word (step 404) from the utterance database 208. Then the ideal duration in seconds for the first word is acquired (step 406) from the utterance database 208. The recognition score for the first word is acquired (step 408) from the utterance database 208. If the individual recognition score for the word is greater than the lower threshold recognition score (60%) then the duration ratio is calculated (step 410) by dividing the actual duration by the ideal duration. If the word is not the last word then the process re-starts at step 404 with the next word in the utterance (step 412). If the word is the last word then an average duration ration is calculated for words with a recognition score above the lower threshold recognition score (step 414). Method 400 ends at step 416.

The tables in Figure 5A and Figure 5B are example utterance sets of words as stored in the utterance database according to a preferred embodiment of the present invention. Referring to Figure 5A there is shown table 500 comprising: recognised words in column A; a recognition score for each word in column B; the actual duration of each word as estimated by the recognition engine 118 in column C; the ideal duration of each word is modeled by the recognition engine 118 in column D; and the duration ratio as calculated by the duration custom server, in column E. Cell B6 of table 500 is the average recognition score calculated by taking an average of the all the individual recognition scores. Cell E6 of table 500 is the average of all the duration ratios with acceptable recognition scores as calculated by the duration custom server.

Figure 5B shows table 502 which has the same column, row and cell definitions as table 500 except that the contents are different. The contents of the cells in Figure 5B correspond to the values of utterance 1 in Figure 6.

Referring to Figure 6 there is shown a table of data returned by the voice recognition engine for three different utterances of the same words "Four three two one". Utterance 1 is spoken slowly, utterance two is spoken normally and utterance 3 is spoken quickly.

Six values are returned for each word in the utterance. For utterance 1 and word "four", the speech recognition engine 118 returns the spelling "4"; the score 50% (which in this case is below the lower threshold recognition score of 60%); the start time 170msec; the end time 870msec; the actual duration (end time-start time=700msec); and the ideal duration 350msec. Similar results are returned for the remaining words in utterance 1, notice that in each case the actual duration of the word is greater than the ideal duration. In the results for utterance 2 notice that the actual duration is of order of the ideal duration for each word. In the results for utterance 3 notice that the actual duration is less than the ideal duration.

Example

John Smith, the user, phones his telephone banking service to query his current account. The phone rings and the IVR answers and plays the first prompt

The IVR prompts "please speak your customer identification number"

The user then utters his four digit id number, "Eight, seven, six, five"

The user's utterance is sent to the speech recognition engine which creates recognition scores and durations for each recognised word (table 500). One word at a time the recognition scores (table 500 column B) are checked, and all found to be above the lower threshold recognition score 60%. Their actual durations (table 500 column C) compared to that of ideal word durations (table 500 column D) to produce individual duration ratio (table 500 column E). The average of the duration ratios is shown in table 500 Cell B6 and is less than one at 0.78 This means that the actual utterance was shorter than the ideal and the user is speaking quicker than the ideal. All the words were successfully recognised (recognition score above 60%) and the average duration ratio is less than the de minimus value

of 0.80 so the application prompts the user to speak more slowly next time. Since the all the words were successfully recognised the application does not re-acquire the utterance at this time.

5 Next, the IVR prompts, "Please speak your password for account eight, seven, six, five and remember to speak a little slower."

The user responds, "four, three, two, one."

10 The user's utterance is sent to the speech recognition engine which creates recognition scores and durations for each recognised word and places them in the utterance database (table 502). One word at a time the recognition score is checked and the score for the first word is found to be lower than the lower threshold recognition score. Ignoring this first
15 word, the duration ratio for the actual and ideal durations for the remaining three words are found (table 502 column E). This time the actual duration for saying the three words is greater than the ideal duration and this indicates that the user is speaking too slowly.

20 From table 502 the duration ratios of the actual duration (Column C) and the ideal duration (column D) for the last three words (column C) is 2.1, 1.9, and 1.7 (column E) which averages 1.9 (Cell E6). Therefore the actual duration for the recognised words is greater than the ideal duration and the user is speaking slower than the ideal. Since the recognition
25 score for the first words is below the lower threshold recognition score (60%) then re-acquisition of the utterance is necessary.

The IVR prompts, "Please speak your password again, remembering to speak slightly slower this time."

30

The user responds with "Four, three, two, one" which hopefully results in a better recognition score and a speed of utterance nearer the ideal.

35

The user's re-acquired utterance is sent to the speech recognition engine which creates recognition scores and durations for each recognised word as before. If the recognition score for the individual words are above the lower threshold recognition score (60%) then the application continues as normal with the rest of the voice application.

40

Although the embodiment has been described in terms of IBM IVR for AIX other IVRs can be used to implement the invention. For instance IBM WebSphere Voice Response for Windows* NT* and Windows 2000 with DirectTalk Technology is an interactive voice response (IVR) product that is for users

who prefer a Windows-based operating environment to run self-service applications. WebSphere Voice Response is capable of supporting simple to complex applications and can scale to thousands of lines in a networked configuration. *Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

The preferred embodiment uses State Table application layer but other embodiments could be implemented in Java or VoiceXML by following the same principles.

Column E, cell B6, cell E6 in tables 500 and 502 are, in this embodiment, only constructs for the purposes of illustration. The duration customer server performs the calculations and stores the values in global or internal variables rather than the utterance database. In another embodiment these values may be stored in the utterance database.

In this embodiment a duration ratio is used to determine the need to prompt for utterance speed and whether to prompt for too fast or too slow. In another embodiment different comparisons between the ideal and the actual duration may be made as long as a magnitude and direction is apparent.

In summary there is disclosed an interactive voice response system, method and computer program product for prompting a user with speech speed feedback during speech recognition. A user who speaks too slowly or too quickly may speak even more slowly or quickly in response to an error in speech recognition. The present system aims to give the user feedback on the speed of speaking. The method comprising: acquiring an utterance from a user; recognising a string of words from the utterance; acquiring for each word the ration of actual duration of delivery to ideal duration; calculating an average ratio for all the words wherein the average ratio is an indication of the speed of the delivery of the utterance; and prompting the user as to the speed of delivery of the utterance according to the average ratio.

CLAIMS

1. An interactive response system comprising:

5 means for acquiring an utterance from a user;

means for recognising a plurality of words from the utterance;

10 means for comparing the actual duration of the utterance with an ideal duration of the recognised words; and

means for prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

15 2. A system as in claim 1 wherein the means for comparing the actual duration of the utterance with an ideal duration of the recognised words comprises:

20 means for acquiring for each word the actual duration of delivery and ideal duration; and

means for comparing the differences in actual duration and ideal duration for each word.

25 3. A system as in claim 2 wherein the means for comparing the difference between actual duration and ideal duration of each word comprises:

30 calculating the ratio of actual duration and ideal duration for each word wherein the ratio is an indication of the speed of delivery of the utterance.

35 4. A system as in claim 3 wherein the means for comparing the actual duration of the utterance with an ideal duration of the recognised words comprises calculating an average of the ratio of words as an indication of the speed of delivery of the utterance.

40 5. A system as in claim 4 wherein each word has an associated recognition score and the ratio of a word is only used in the average calculation if the associated recognition score is above a lower threshold recognition score.

6. A system as in any of claims 1 to 5 wherein if the actual duration is greater than the ideal duration then prompting the user that the utterance has been delivered too fast and if the actual duration is less than the

ideal duration then prompting the user that the utterance has been delivered too slow.

5 7. A system as in any of claims 1 to 6 wherein the means for prompting the user only operates if the actual duration and the ideal duration differ by more than a de minimus value.

8. A method in an interactive response system comprising:

10 acquiring an utterance from a user;

recognising a plurality of words from the utterance;

15 comparing the actual duration of the utterance with an ideal duration of the recognised words; and

prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

20 9. A computer program product for processing one or more sets of data processing tasks, said computer program product comprising computer program instructions stored on a computer-readable storage medium for, when loaded into a computer and executed, causing a computer to carry out the steps of: acquiring an utterance from a user;

25 recognising a plurality of words from the utterance;

comparing the actual duration of the utterance with an ideal duration of the recognised words; and

30 prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

ABSTRACT**METHOD AND APPARATUS FOR A INTERACTIVE VOICE RESPONSE SYSTEM**

5 There is disclosed an interactive voice response system for prompting
a user with feedback during speech recognition. A user who speaks too
slowly or too quickly may speak even more slowly or quickly in response to
an error in speech recognition. The present system aims to give the user
specific feedback on the speed of speaking. The method comprising:
10 acquiring an utterance from a user; recognising a string of words from the
utterance; acquiring for each word the ration of actual duration of
delivery to ideal duration; calculating an average ratio for all the words
wherein the average ratio is an indication of the speed of the delivery of
the utterance; and prompting the user as to the speed of delivery of the
15 utterance according to the average ratio.

